

# Effects of noise on convergent game learning dynamics

**James BT Sanders<sup>1</sup>, Tobias Galla<sup>1</sup> and Jonathan L Shapiro<sup>2,3</sup>**

<sup>1</sup>Theoretical Physics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, United Kingdom

<sup>2</sup>School of Computer Science, The University of Manchester, Manchester M13 9PL, United Kingdom

<sup>3</sup>Centre for Interdisciplinary Computational and Dynamical Analysis (CICADA), The University of Manchester, Manchester M13 9PL, United Kingdom

**Abstract.** We study stochastic effects on the lagging anchor dynamics, a reinforcement learning algorithm used to learn successful strategies in iterated games, which is known to converge to Nash points in the absence of noise. The dynamics is stochastic when players only have limited information about their opponents' strategic propensities. The effects of this noise are studied analytically in the case where it is small but finite, and we show that the statistics and correlation properties of fluctuations can be computed to a high accuracy. We find that the system can exhibit quasicycles, driven by intrinsic noise. If players are asymmetric and use different parameters for their learning, a net payoff advantage can be achieved due to these stochastic oscillations around the deterministic equilibrium.

E-mail: James.Sanders-2@postgrad.manchester.ac.uk,  
Tobias.Galla@manchester.ac.uk, jls@cs.man.ac.uk

PACS numbers: 02.50.Le, 05.10.Gg, 02.50.Ey

## 1. Introduction

Competitive situations in biology, the social sciences, and computer science are often modelled using the formalism of game theory [1]. In these models, two or more players each choose between a set of actions, and each player then receives a reward (or ‘payoff’), depending on their own action and those of the other players. Traditionally, games are classified in terms of outcomes one would expect for perfectly rational players, so-called Nash equilibria [2, 3]. These are points in strategy space such that no player can improve their expected payoff by unilaterally changing their own behaviour. In many simple games, it is straightforward to compute such equilibrium points. Typical equilibrium strategies will be probabilistic—these are known as ‘mixed strategies’, while a deterministic strategy is described as ‘pure’. In games with a large number of actions to choose between, there may be many equilibria, or even none at all if there are infinitely many actions. Even when a unique optimal solution to a game is available, the players

may not have enough information or computational capacity to identify it. To model situations where an optimal strategy does not exist, or where the players are unable to adopt it, learning algorithms have been introduced. These describe scenarios in which players repeatedly play a game, modifying their strategies each time to try and maximize the payoffs they receive [4, 5, 6, 7, 8, 9, 10]. The purposes of these algorithms vary. Some are attempts to model natural or social systems. Other learning algorithms are used in decision-making software, optimization, and more generally in machine learning, a branch of computer science concerned with developing algorithms to improve behaviour using empirical information [11]. In our work we do not aim to model the psychological processes underpinning human learning, but instead focus on a class of machine-learning algorithms proposed by Dahl [12], the so-called ‘lagging anchor dynamics’. Before we define the exact details of this dynamics it is useful to briefly outline some of the basic general principles of game learning.

Learning algorithms typically require a player to play a mixed strategy, iteratively modifying the probabilities of playing each action depending on the outcomes of past games. In general, the player gradually increases the probability of playing actions that would achieve higher payoffs against their opponent’s current strategy, while decreasing the probability of those that would perform poorly [10, 13]. One well-studied adaptation mechanism in machine learning is reinforcement learning. In this class of dynamics, the only information available to a player is the relative success or failure of previous actions [11, 14].

In game learning the expected payoff can be viewed as a function of the players’ mixed strategies, and each player wishes to maximize their own payoff, so that reinforcement learning can be seen as an optimization problem. One well-known optimization method is gradient ascent, an iterative scheme in which the maximum of a function is approached by repeatedly taking small steps in the direction in which the value of the function increases most quickly. If  $\mathbf{x}^{(0)}$  is an initial estimate of the maximum of the function  $f(\cdot)$ , then gradient ascent proceeds according to the recurrence relation

$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} + \kappa \nabla f(\mathbf{x}^{(n)}), \quad (1)$$

where the step size  $\kappa$  is a small positive constant [15]. In the context of games, each player views their expected payoff as a function of their mixed strategy, imagining the other players’ strategies are constant, and carries out one step of the gradient ascent method. This is sometimes known as ‘simultaneous gradient ascent’ [16]. If each player has complete knowledge of their opponents’ mixed strategies, gradient ascent can be used deterministically. If not, the players must estimate their opponents’ strategies based on observations. This introduces noise into the system, as the opponents’ actions are drawn probabilistically from an underlying mixed strategy, which is unknown to the observer.

In a game with a single Nash equilibrium point, it is often argued that rational players would choose the equilibrium strategies, in which case, it is desirable that a learning algorithm should find these strategies. The standard simultaneous gradient

ascent algorithm is typically able to converge to pure equilibrium strategies, but not mixed strategies [12]. A number of variations of the algorithm have been proposed to address this issue. For example, in some versions the players attempt to predict their opponents' future strategies, and calculate the gradient of the payoff function based on these expected future strategies rather than the current ones. Another variation is 'satisficing', in which each player maintains an 'aspiration', and only updates their strategy if their expected payoff is lower than their current aspiration [12].

The lagging anchor algorithm is another modification of gradient ascent, in which each player maintains a long-term memory of past strategies—a 'lagging anchor'—which is coupled to their current strategy. This algorithm was introduced by Dahl [12], who concentrated on its deterministic behaviour, when each player knows precisely their opponents' current mixed strategies. He was able to prove that the dynamics is able to converge to mixed equilibrium strategies in a broad class of games. Butterworth and Shapiro [17] looked at the continuous-time limit of the system, deriving further results concerning the deterministic behaviour. They also investigated two different situations in which the players have limited information, noting the appearance of stochastic quasicycles.

The main purpose of our work is to extend the analysis of stochastic lagging anchor learning. We concentrate on situations where the intrinsic noise is small but finite, and examine analytically its effect upon the system. In particular, we investigate how stochasticity affects the stability of the dynamics, and whether players can exploit noise in the system to increase their own payoffs.

## 2. Model definitions

### 2.1. General definitions

Consider a finite, two-player game with  $m$  actions available to player 1, and  $n$  for player 2. A particular game is defined by two  $m \times n$  payoff matrices,  $E^{(1)}$  and  $E^{(2)}$ , so that if player 1 chooses action  $\alpha$ , and player 2 chooses action  $\beta$ , the payoff to player  $i$  is  $E_{\alpha\beta}^{(i)}$ .

At any given time step  $t$ , the players have mixed strategies represented by two vectors  $\mathbf{p}(t)$  and  $\mathbf{q}(t)$ , and one instance of the game is played. Player 1 chooses action  $\alpha$  with probability  $p_\alpha(t)$ , while player 2 chooses action  $\beta$  with probability  $q_\beta(t)$ . The corresponding payoffs are distributed to the players, who then update their strategies according to their learning method (this will be described below).

For simplicity, we initially focus on the so-called 'matching pennies' game [1]. In this zero-sum game, each player has a coin and selects 'heads' or 'tails'. If the players' choices match, player 1 keeps both pennies, otherwise player 2 keeps them both. The payoff matrices for this two-action game are

$$E^{(1)} = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} = -E^{(2)}.$$

The unique Nash equilibrium for this game is the pair of mixed strategies for which each player chooses heads or tails with equal probabilities. For convenience, we write the mixed strategies as  $\mathbf{p}(t) = (\frac{1}{2} + x(t), \frac{1}{2} - x(t))^T$  and  $\mathbf{q}(t) = (\frac{1}{2} + y(t), \frac{1}{2} - y(t))^T$ , so that the Nash equilibrium is at  $x = y = 0$ , the expected payoff to player 1 is  $\mathbf{p}^T E^{(1)} \mathbf{q} = 4xy$ , and the expected payoff to player 2 is  $-4xy$ .

## 2.2. Gradient ascent

Rephrasing the matching pennies game as an optimization problem, player 1 wishes to maximize  $4xy$ , and is able to vary  $x$  in the interval  $[-\frac{1}{2}, \frac{1}{2}]$ . Player 2 wants to maximize  $-4xy$ , and can vary  $y$  over the same interval. If each player applies the gradient ascent method (1), the update rules read

$$\begin{aligned} x(t+1) &= x(t) + 4\kappa_1 \tilde{y}(t), \\ y(t+1) &= y(t) - 4\kappa_2 \tilde{x}(t), \end{aligned} \tag{2}$$

where  $\kappa_1$  and  $\kappa_2$  are the two players' gradient ascent step sizes, and can be thought of as learning rates. In the event that  $x$  or  $y$  are taken out of their allowed intervals by these update rules, they are simply mapped back to the nearest allowed points. Here we use the notation  $\tilde{x}(t)$  to represent player 1's estimate of player 2's strategy at time  $t$ , and respectively for  $\tilde{y}(t)$ . We assume the players calculate these estimates using a geometric discounted average of their opponents' previous choices as in [17], giving update rules for  $\tilde{x}$  and  $\tilde{y}$ ,

$$\begin{aligned} \tilde{x}(t+1) &= \tilde{x}(t) + \phi_1 (X(t) - \tilde{x}(t)), \\ \tilde{y}(t+1) &= \tilde{y}(t) + \phi_2 (Y(t) - \tilde{y}(t)), \end{aligned}$$

where  $\phi_1$  and  $\phi_2$  are constant parameters, and  $X(t)$  and  $Y(t)$  represent the pure strategies chosen by players 1 and 2, respectively, at time  $t$ , taking the value  $+\frac{1}{2}$  for action 1, and  $-\frac{1}{2}$  for action 2. Deterministic mean field equations are recovered by setting  $X = x$  and  $Y = y$ .<sup>‡§</sup>

## 2.3. Lagging anchor dynamics

Although the Nash equilibrium ( $x = y = 0$ ) is a fixed point of the deterministic gradient ascent equations (2), it is not stable and the system will typically not converge to this equilibrium. This is a general feature of mixed Nash equilibria under gradient-type strategy updates, as has been known for a long time (see, for example [18, 19]).

<sup>‡</sup> Another method suggested in [17] covers situations where the players cannot observe their opponents' actions, and must calculate gradient estimates using only observations of their own payoffs. Using this method instead of opponent modelling appears to have little effect on the dynamics, except to increase the size of stochastic effects.

<sup>§</sup> Here, we follow Butterworth and Shapiro in updating the strategies at time  $t$  based on the estimates from time  $t-1$ , which only include contributions from observations from time  $t-2$  and earlier. It could be argued that the players should include observations from time step  $t-1$  in their strategy update at time  $t$ , but this would be unlikely to change the dynamics significantly.

Several mechanisms have been proposed to make equilibria stable under gradient-type updates, among them the lagging anchor algorithm we study in this work. The lagging anchor algorithm involves introducing an additional state variable for each player. These variables, called ‘lagging anchors’, are weighted averages of previously used mixed strategies. For matching pennies, the lagging anchor dynamics, with the same geometric discounted estimates, is

$$\begin{aligned} x(t+1) &= x(t) + 4\kappa_1 \tilde{y}(t) + \mu_1 (\bar{x}(t) - x(t)), \\ y(t+1) &= y(t) - 4\kappa_2 \tilde{x}(t) + \mu_2 (\bar{y}(t) - y(t)), \\ \bar{x}(t+1) &= \bar{x}(t) + \nu_1 (x(t) - \bar{x}(t)), \\ \bar{y}(t+1) &= \bar{y}(t) + \nu_2 (y(t) - \bar{y}(t)), \\ \tilde{x}(t+1) &= \tilde{x}(t) + \phi_1 (X(t) - \tilde{x}(t)), \\ \tilde{y}(t+1) &= \tilde{y}(t) + \phi_2 (Y(t) - \tilde{y}(t)), \end{aligned} \tag{3}$$

where the lagging anchors are given by  $\bar{\mathbf{p}}(t) = (\frac{1}{2} + \bar{x}(t), \frac{1}{2} - \bar{x}(t))^T$  and  $\bar{\mathbf{q}}(t) = (\frac{1}{2} + \bar{y}(t), \frac{1}{2} - \bar{y}(t))^T$ . The introduction of the anchors does not change the location of the equilibrium of the deterministic dynamics, but it does stabilize it if the parameters are chosen appropriately. The variables  $\mu_i$  and  $\nu_i$  ( $i = 1, 2$ ) are non-negative model parameters, whose interpretation will be detailed further below. The system of equations (3) is a generalization of that presented in [17], which is recovered if  $\kappa_1 = \kappa_2$ ,  $\mu_1 = \mu_2 = \nu_1 = \nu_2$ , and  $\phi_1 = \phi_2$ . As before we will implicitly assume that any variable leaving the interval  $[-1/2, 1/2]$  will be clipped and mapped onto the points  $1/2$  or  $-1/2$  respectively. This introduces an occasional nonlinearity into a dynamics whose deterministic behaviour would otherwise be linear.

The update rules (3) contain eight constant parameters. The gradient ascent step sizes  $\kappa_1$  and  $\kappa_2$  can be thought of as learning rates, and must be small and positive for the dynamics to be capable of learning the equilibrium strategy. The anchor parameters  $\mu_i$  and  $\nu_i$  determine the strength of coupling between the strategies and anchors. A large  $\nu_i$  corresponds to anchors that are pulled strongly towards the mixed strategies, and a large  $\mu_i$  corresponds to strategies that are pulled strongly towards the anchors. In particular, setting  $\mu_i$  to zero or  $\nu_i$  to one is equivalent to removing the lagging anchors altogether, and setting  $\nu_i$  to zero will hold the anchors in their initial positions, permanently biasing the strategies towards these points. Finally, the modelling rates  $\phi_1$  and  $\phi_2$  represent the rates at which the players update their models of their opponents’ strategies. A value of  $\phi_i = 1$  is equivalent to the players discarding all previous information about their opponents’ actions and assuming that their strategy is simply to play their most recently used action every game. A value of  $\phi_i$  between 0 and 1 will cause the opponent models to be averages of past behaviour, weighted according to a geometric discounting factor.

#### 2.4. Compact notation

For convenience, we rewrite the full stochastic lagging anchor system (3) in matrix form, separating the deterministic and stochastic effects. Since the observed pure strategies

$X$  and  $Y$  are unbiased estimates of the mixed strategies  $x$  and  $y$ , respectively, we can write them in terms of mean-zero noise terms  $\xi$  and  $\chi$ ,

$$\begin{aligned} X(t) &= x(t) + \xi(t), \\ Y(t) &= y(t) + \chi(t). \end{aligned} \tag{4}$$

More precisely,  $\xi(t)$  takes the value  $1/2 - x(t)$  with probability  $1/2 + x(t)$ , and the value  $-1/2 - x(t)$  with probability  $1/2 - x(t)$ , and similar for  $\chi(t)$ . The state of the system at any given time is fully described by the vector  $\zeta = (x \ y \ \bar{x} \ \bar{y} \ \tilde{x} \ \tilde{y})^T$ . The update rules (3) can be written in matrix form as

$$\zeta(t+1) = J\zeta(t) + \varphi(t), \tag{5}$$

where the column vector  $\varphi = (0 \ 0 \ 0 \ 0 \ \phi_1 \xi \ \phi_2 \chi)^T$  contains the noise terms, and the constant matrix  $J$  is

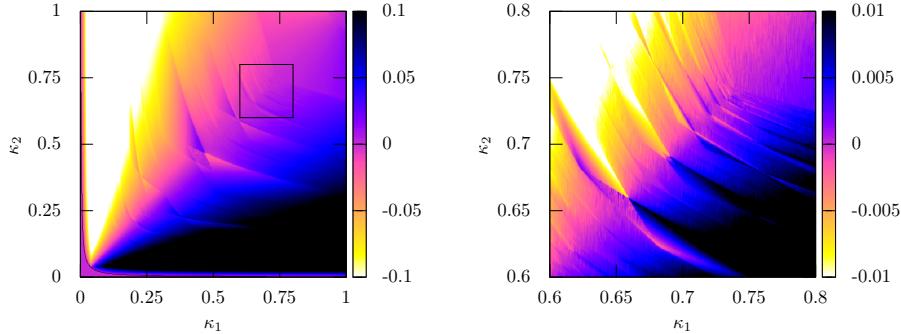
$$J = \begin{pmatrix} 1 - \mu_1 & 0 & \mu_1 & 0 & 0 & 4\kappa_1 \\ 0 & 1 - \mu_2 & 0 & \mu_2 & -4\kappa_2 & 0 \\ \nu_1 & 0 & 1 - \nu_1 & 0 & 0 & 0 \\ 0 & \nu_2 & 0 & 1 - \nu_2 & 0 & 0 \\ \phi_1 & 0 & 0 & 0 & 1 - \phi_1 & 0 \\ 0 & \phi_2 & 0 & 0 & 0 & 1 - \phi_2 \end{pmatrix}.$$

### 3. Analytical analysis

In this section we will investigate the lagging anchor dynamics on a theoretical level. We first briefly discuss the outcome of deterministic learning, and then derive analytical expressions characterizing the properties of stochastic learning in the limit of small noise.

#### 3.1. Deterministic system

The deterministic limit of the lagging anchor update is obtained by replacing  $X(t) \rightarrow x(t)$  and  $Y(t) \rightarrow y(t)$  in (3), or equivalently by setting  $\xi(t) = \chi(t) = 0$ . The resulting dynamics have been studied previously by Dahl [20, 21, 12] and by Butterworth and Shapiro in [17]. Dahl considered the case where  $\kappa_1 = \kappa_2$ ,  $\mu_1 = \mu_2 = \nu_1 = \nu_2$ , and the game is zero-sum, and showed that, provided the learning rate and anchor parameter are small enough and the payoff matrix is invertible, a broad class of mixed equilibrium strategies are asymptotically stable. The authors of [17] continued this analysis, and determined the region of the parameter space for which the dynamics is stable. We do not repeat the full analysis for the discrete-time system here, as it is straightforward to see that the stability of the Nash point depends on the eigenvalues of the matrix  $J$  in (5). In the interior of the strategy simplexes, the deterministic dynamics is simply given by the linear map  $\zeta(t+1) = J\zeta(t)$ , so that  $\zeta = 0$  is a stable fixed point provided all eigenvalues of  $J$  have moduli less than one. In this case, the strategies will converge to the equilibrium point, no matter from what initial condition the dynamics is started. Asymptotically the payoffs to the two players will be those at the equilibrium point, i.e. the long-term payoff for each player vanishes in the matching-pennies game.



**Figure 1.** (Colour on-line) Colour maps of the average payoff to player 1 for various values of  $\kappa_1$  and  $\kappa_2$ . The right plot covers the region marked by a square in the left plot, and shows clear fractal behaviour.

When the matrix  $J$  has one or more eigenvalues outside the unit circle of the complex plane, the strategies spiral outwards until they reach the simplex boundaries. Then the nonlinearity sets in, effectively clipping the phase space of the system. Numerically iterating the deterministic map we typically observe periodic motion. Interestingly, we find that if the players are not symmetric (i.e., if they have different learning rates, anchor parameters, or modelling rates), it is possible for one player to consistently achieve a positive payoff. The dependence of this average payoff on the parameters appears to be quite complex, and shows discontinuities in parameter space, as well as fractal-like behaviour, as shown in figure 1.

### 3.2. Stochastic system

The authors of [17] determined the form of the covariance of the players' strategies in the continuous limit of the stochastic dynamics, concluding that where the deterministic dynamics is stable, the strategies will typically perform quasicycles about the equilibrium point, with a frequency determined by the eigenvalues of the Jacobian matrix. We continue this analysis by calculating analytic approximations of the covariances and power spectra of the state variables, and comparing them to results from simulations.

**3.2.1. Assumptions** The chief barrier to analytically calculating statistics of the lagging anchor system (3) is the presence of the truncations that are applied if the strategies leave their allowed intervals. We assume that the parameters of the algorithm are such that the deterministic dynamics converges, and that the intrinsic noise caused by the limited information is small. Under these conditions, the strategies spend most of their time near the equilibrium point and so the probability of them reaching the boundaries of the simplexes is negligible, and the nonlinear corrections can be ignored.

*3.2.2. Correlations of the noise terms* The correlation matrix  $D(t) = \langle \varphi(t)\varphi(t)^T \rangle$  can be written in block form as

$$D(t) = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \phi_1^2 \langle \xi(t)^2 \rangle & 0 \\ 0 & 0 & \phi_2^2 \langle \chi(t)^2 \rangle \end{pmatrix},$$

since the two noise terms  $\xi(t)$  and  $\chi(t)$  are independent. The variance of the noise terms  $\xi$  and  $\chi$  can be calculated from their definitions (4), leading to

$$\langle \xi^2 \rangle = \frac{1}{4} - x^2, \quad \langle \chi^2 \rangle = \frac{1}{4} - y^2.$$

To simplify the calculation of the covariances and power spectra of the strategies, we neglect the quadratic terms in  $x$  and  $y$  in these expressions. This assumption can be expected to be accurate close to the Nash point.<sup>||</sup> Under this simplifying assumption,  $D(t)$  can be regarded as constant,

$$D(t) = D = \frac{1}{4} \begin{pmatrix} 0 & 0 & 0 \\ 0 & \phi_1^2 & 0 \\ 0 & 0 & \phi_2^2 \end{pmatrix}.$$

*3.2.3. Correlations of the state variables* We calculate an approximation of the equal-time covariance matrix  $C(t)$  of the state vector  $\zeta(t)$ ,

$$C(t) = \langle \zeta(t)\zeta(t)^T \rangle.$$

Using the matrix form (5) of the lagging anchor algorithm and the fact that  $\varphi(t)$  and  $\zeta(t)$  are uncorrelated, we have

$$C(t+1) = JC(t)J^T + D. \tag{6}$$

We are interested in the long-term value of  $C(t)$ , after transients have died away. In this asymptotic regime the components of  $C$  are related to the size of the stochastic deviations from the equilibrium point. If this limiting value exists, it is given by the fixed point  $C^*$  of (6),

$$C^* = JC^*J^T + D. \tag{7}$$

This matrix equation is easily solved by rearranging it into a linear system [26]. It has a unique positive definite solution whenever the deterministic lagging anchor dynamics has a stable fixed point [27].

The long-term expected payoff  $\langle u \rangle$  received per game by player 1 can be calculated given  $C^*$ , using

$$\langle u \rangle = 4 \langle xy \rangle = 4C_{12}^*.$$

<sup>||</sup> Similar to [22, 23, 24, 25] this is equivalent to a systematic expansion in the amplitude of the intrinsic noise.

*3.2.4. Power spectra* To check whether the stochastic dynamics display quasicycles, we define a power spectral density matrix for the vector  $\zeta$ ,

$$P(\omega) = \lim_{L \rightarrow \infty} \frac{1}{L} \langle \hat{\zeta}(\omega, L) \hat{\zeta}(\omega, L)^\dagger \rangle,$$

where  $\dagger$  is the Hermitian conjugate, and the hats denote discrete Fourier transforms,

$$\hat{\zeta}(\omega, L) = \sum_{t=0}^{L-1} \zeta(t) \exp(-i\omega t).$$

An expression for the power spectra can be obtained by taking the Fourier transform of the lagging anchor recurrence relation (5), which, for  $L \gg 1$ , transforms the system into

$$(e^{i\omega} I - J) \hat{\zeta}(\omega, L) = \hat{\varphi}(\omega, L),$$

where  $I$  is the  $6 \times 6$  identity matrix. Letting  $M(\omega) = e^{i\omega} I - J$  we obtain

$$P(\omega) = \lim_{L \rightarrow \infty} \frac{1}{L} M(\omega)^{-1} \langle \hat{\varphi}(\omega, L) \hat{\varphi}(\omega, L)^\dagger \rangle M(\omega)^{\dagger -1}. \quad (8)$$

If we write the expectation value in this last expression in terms of the noise terms,

$$\langle \hat{\varphi}(\omega, L) \hat{\varphi}(\omega, L)^\dagger \rangle = \sum_{t=0}^{L-1} \sum_{s=0}^{L-1} \langle \varphi(t) \varphi(s)^T \rangle e^{-i\omega(t-s)}.$$

Using the fact

$$\langle \varphi(t) \varphi(s)^T \rangle = D \delta_{ts},$$

we find

$$\langle \hat{\varphi}(\omega, L) \hat{\varphi}(\omega, L)^\dagger \rangle = LD.$$

Substituting this into (8) gives the result

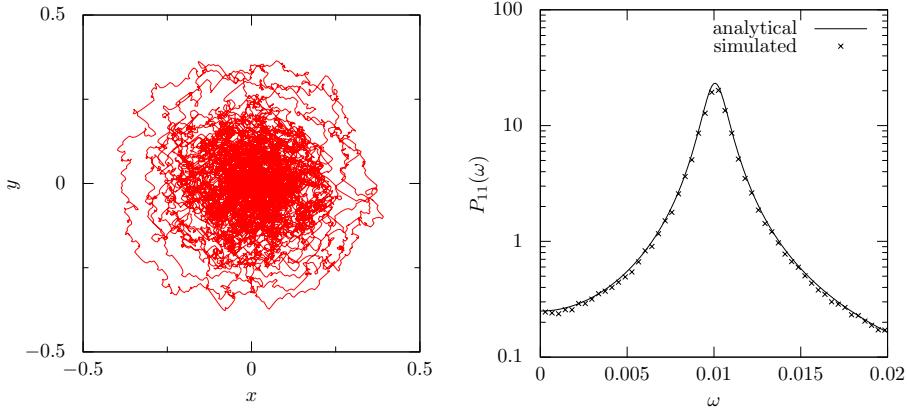
$$P(\omega) = M(\omega)^{-1} DM(\omega)^{\dagger -1}. \quad (9)$$

This expression is similar to those obtained for population-based models by means of the celebrated van Kampen expansion, see e.g. [28, 29, 30]. Studies of intrinsic noise in game learning can also be found in [22, 23, 24, 25]. Expression (9) provides an explicit prediction for the spectra properties of fluctuations about the Nash point (or equivalently a prediction of its correlation function). These predictions can be tested against simulations, using a large finite value of  $L$ .

## 4. Test against simulations

### 4.1. Identical players

First, we consider the case where the two players have identical update rules (i.e.,  $\kappa_1 = \kappa_2$ ,  $\mu_1 = \mu_2$ ,  $\nu_1 = \nu_2$ , and  $\phi_1 = \phi_2$ ). We run simulations of the lagging anchor algorithm for the matching pennies game, comparing the results to the analytic calculations discussed above, for various values of the four parameters.



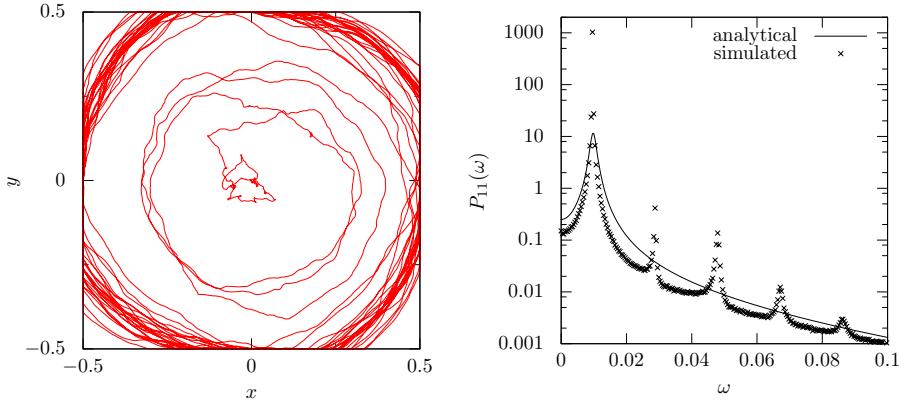
**Figure 2.** (Colour on-line) **Left:** Phase plot of the strategy coordinates  $x$  and  $y$  of the two players in a simulation of the lagging anchor algorithm, for the matching pennies game. The parameters have the values  $\kappa_i = 0.005$ ,  $\mu_i = \nu_i = 0.05$ ,  $\phi_i = 0.5$ , for which the deterministic dynamics converges to the equilibrium point. **Right:** Corresponding power spectrum of the first player’s strategy coordinate  $x$ . The solid curve shows the spectrum calculated analytically using (9), while the crosses show a simulated spectrum, averaged over 1000 realizations of the dynamics ( $L = 2^{14}$ ).

The left-hand panel of figure 2 shows a phase plot of the players’ strategies in one realization, for values of the parameters where the deterministic dynamics is stable. The right-hand panel shows the associated spectral density of player one’s strategy, which shows a single sharp peak, reflecting the existence of quasicycles. The analytic approximation of the power spectrum is very close to the simulated spectrum, suggesting that the assumptions used to calculate it—i.e., ignoring the restriction of  $x$  and  $y$  to intervals, and neglecting the quadratic terms in  $D(t)$ —hold for these values of the parameters.

In contrast, figure 3 (left-hand panel) shows a phase plot of simulations for another set of parameters, where the deterministic dynamics is unstable. The strategies perform a noisy version of the limit cycle found in the deterministic dynamics, generally staying near the boundaries. The power spectrum of player 1’s strategy, shown in the right-hand panel is not in good agreement with the analytic calculation (9) as the strategies do not remain close to the equilibrium point. However, the analytic spectrum correctly predicts the location of the main peak.

Perhaps the most important diagnostic of the effects of changing the parameters is the long-term average variance of the components of the strategies. This quantifies the size of stochastic oscillations about the equilibrium point, and therefore the success of the algorithm in learning the equilibrium strategies.

Figure 4 shows the dependence of the size of the oscillations on the anchor parameters  $\mu$  and  $\nu$ . In general, it appears that in order to have small oscillations, the memory-loss parameter  $\nu$  must be small compared to  $\mu$ . In other words, the lagging anchors must move slowly and pull the strategies towards them strongly. The opposite

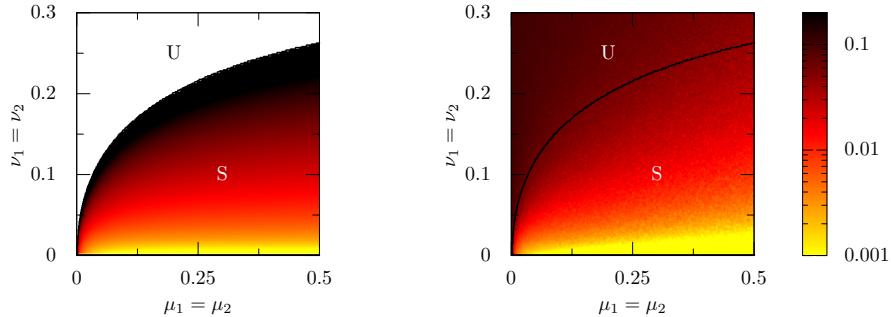


**Figure 3.** (Colour on-line) **Left:** Phase plot of the strategy coordinates  $x$  and  $y$  of the two players in a simulation of the lagging anchor algorithm, for the matching pennies game. The parameters have the values  $\kappa_i = 0.005$ ,  $\mu_i = \nu_i = 0.05$ ,  $\phi_i = 0.05$ , for which the deterministic dynamics do not converge. **Right:** Corresponding power spectrum of the first player’s strategy coordinate  $x$ . The solid curve shows the values calculated analytically using (9), while the crosses show a simulated spectrum, averaged over 1000 realizations of the dynamics. The sharp peaks in the simulated spectrum are reminiscent of those of a noisy square wave, demonstrating the dominance of the nonlinear truncation of the strategies for these parameters.

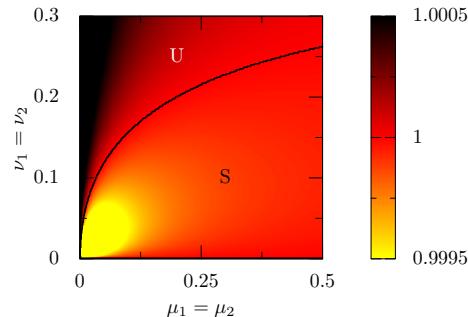
situation, in which the anchors quickly move towards the strategies, but only weakly affect them, is comparable to removing the anchor terms.

Far inside the region where the deterministic dynamics is stable, oscillations are found to be small, and the analytic and simulated variances are very similar. As the stability line is approached, the analytic approximation of the variance diverges, as the simplex boundaries were ignored in the calculation. The variance measured in simulations approaches a constant value, as the strategies in the simulations are not allowed to leave the bounded intervals on which they are defined.

The long-term variance of the strategies does not completely characterize the behaviour of the algorithm, however. In the low- $\nu$  region of the plots in figure 4, the anchors will move very slowly, so that if they are not initially close to the equilibrium point, it will take a long time for transients to die away. The colour map in figure 5 shows the largest modulus of the eigenvalues of  $J$ , which we denote by  $\lambda$ . This quantity determines the stability of the deterministic dynamics. Where  $\lambda$  is significantly smaller than one, the deterministic dynamics quickly converge, and so transients in the stochastic dynamics are short-lived. Where  $\lambda$  is close to, but smaller than one, the deterministic dynamics converges slowly, so the transients are long-lasting. If the lagging anchor algorithm were used in a practical situation, it may be necessary to choose the anchor parameters to strike a balance between having small oscillations and short-lived transients.



**Figure 4.** (Colour on-line) Colour maps showing the dependence of the variance of player 1's strategy coordinate  $x$  on the anchor parameters  $\mu$  and  $\nu$ , while keeping the other parameters fixed at  $\kappa_i = 0.005$ ,  $\phi_i = 0.5$ , for the matching pennies game. The panel on the left shows the analytic approximation, and the right-hand panel the results from simulations (averaged over multiple runs), with the strategies and anchors beginning at the equilibrium point. In each case, the solid line shows where the deterministic dynamics become unstable. In the left-hand panel, the variance diverges as the instability region  $U$  is approached.

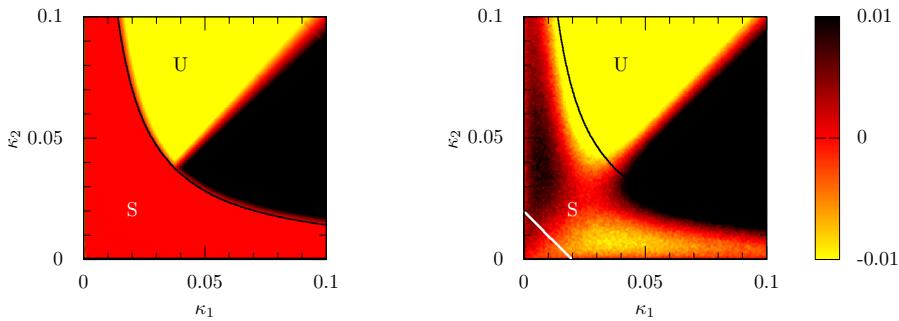


**Figure 5.** (Colour on-line) Colour map showing the dependence of the largest eigenvalue modulus  $\lambda$  of  $J$  on the anchor parameters  $\mu$  and  $\nu$  for the same values of  $\kappa$  and  $\phi$  as in figure 4. The solid black curve is the level set for which this modulus is one, and marks the boundary between the stable (S) and unstable (U) behaviours of the deterministic dynamics.

#### 4.2. Two-player learning with non-identical players

We now turn to situations in which the two players use different values of the parameters  $\kappa$ ,  $\mu$ ,  $\nu$ , and  $\phi$ , still using the matching pennies game. We keep all but one of the parameters constant, then run simulations for a range of values of the remaining parameter for each player.

An interesting quantity to consider is the long-term average payoff to one of the players (recall that we are considering a zero-sum game in which the players receive equal

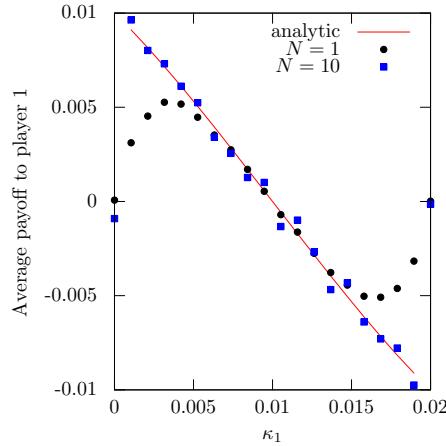


**Figure 6.** (Colour on-line) Colour maps showing the average payoff to player 1 for varying values of  $\kappa_1$  and  $\kappa_2$ , with the other parameters kept fixed at  $\mu_i = \nu_i = 0.05$  and  $\phi_i = 0.5$ . Left: deterministic dynamics, right: stochastic learning. The solid black lines mark the boundary between the stable (S) and unstable (U) phases of the deterministic dynamics. Figure 7 compares the stochastic behaviour along the solid white line  $\kappa_1 + \kappa_2 = 0.02$  with analytic predictions.

and opposite payoffs). Figure 6 shows the dependence of this quantity on the learning rates of the two players, while the other parameters are kept constant, in both the stochastic and deterministic dynamics. Outside the stability region, the deterministic and stochastic behaviour are similar—the player with the highest learning rate wins. This is due to the dominant influence of the truncation of the strategies to their intervals. The structures do not match exactly however—while a player’s payoff reduces to zero as the stability region is approached in the deterministic case, it remains nonzero a little way beyond the stability line in the stochastic dynamics—this appears to be because, just inside the stability region, the stochastic quasicycles push the strategies outwards to the boundaries, where they mimic the limit cycles, and payoffs, found in the unstable deterministic case.

Where the deterministic dynamics is stable, it converges to the equilibrium point, where the average payoff is zero, but this behaviour is not seen in the stochastic dynamics, in which a player can still obtain a finite payoff well inside the stability region. Similar effects of increased payoffs to one player driven by intrinsic noise have been observed in learning processes based on noisy replicator dynamics in [31, 32].

Figure 7 compares the analytic values of the average payoff far inside the stability region with numerical simulations. The predictions are accurate near the  $\kappa_1 = \kappa_2$  line, but systematic deviations become apparent as  $\kappa_1$  or  $\kappa_2$  approach zero, as one player’s strategies are able to move far away from the equilibrium, so the small-noise approximation becomes inappropriate. In fact adopting a batch learning approach (see [22, 23]) increases the range of accuracy of the theory. In batch learning players play ‘batches’ of  $N$  observations at each update step, i.e. the random variables  $X(t)$  and  $Y(t)$  in (3) are replaced by averages over  $N$  independent draws of actions from the players’ mixed strategies. This procedure reduces the noise level (which scales as



**Figure 7.** (Colour on-line) Plot showing the payoff to player 1 as the learning rates are varied along the line  $\kappa_1 + \kappa_2 = 0.02$ , with other parameters as in figure 6. The solid line shows the analytic values, squares the simulated values. Circles show how the simulated values change if there is batching—the players only update their strategies every  $N$  games, where in this case  $N = 10$ . Batching reduces the size of stochastic effects, improving the match with the analytic calculations.

$N^{-1/2}$ ), and as seen in figure 7 the analytic predictions become more accurate as  $N$  is increased.

#### 4.3. Other games

The lagging anchor algorithm can easily be applied to other games. Suppose a game has  $m$  pure strategies for player 1 and  $n$  for player two. Let player 1's strategy be represented by the vector  $\mathbf{p}$ , and player 2's by  $\mathbf{q}$ , and let  $\mathbf{p}^*$  and  $\mathbf{q}^*$  be the mixed strategies at an equilibrium point. We introduce strategy coordinates  $\mathbf{x}$  and  $\mathbf{y}$  defined by

$$\begin{aligned}\mathbf{p} &= L^{(1)}\mathbf{x} + \mathbf{p}^*, \\ \mathbf{q} &= L^{(2)}\mathbf{y} + \mathbf{q}^*,\end{aligned}$$

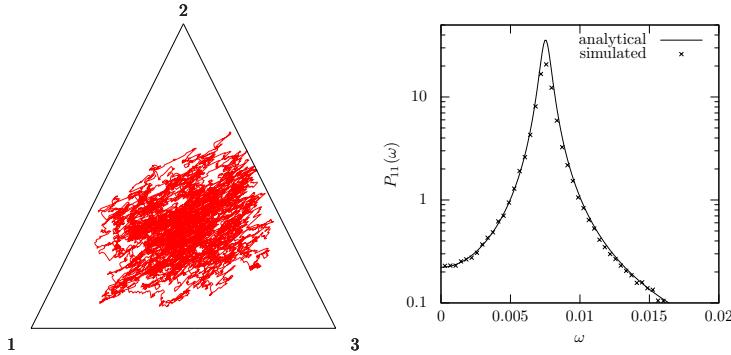
where  $L^{(1)}$  and  $L^{(2)}$  are  $m \times (m-1)$  and  $n \times (n-1)$  matrices, respectively, of the form

$$L^{(1)} = \begin{pmatrix} I_{m-1} \\ -\mathbf{1}_{m-1}^T \end{pmatrix}, \quad L^{(2)} = \begin{pmatrix} I_{n-1} \\ -\mathbf{1}_{n-1}^T \end{pmatrix},$$

where  $I_k$  denotes an identity matrix of size  $k \times k$ , and  $\mathbf{1}_k$  a column vector of ones of length  $k$ . Let the players' payoff matrices be  $E^{(1)}$  and  $E^{(2)}$ , respectively. Then the expected payoffs  $u_1$  and  $u_2$  can be written

$$\begin{aligned}u_1 &= \mathbf{x}^T A^{(1)} \mathbf{y} + \mathbf{p}^{*T} E^{(1)} L^{(2)} \mathbf{y} + \mathbf{p}^{*T} E^{(1)} \mathbf{q}^*, \\ u_2 &= \mathbf{x}^T A^{(2)} \mathbf{y} + \mathbf{q}^{*T} E^{(2)^T} L^{(1)} \mathbf{x} + \mathbf{p}^{*T} E^{(2)} \mathbf{q}^*,\end{aligned}\tag{10}$$

where  $A^{(1)} = L^{(1)^T} E^{(1)} L^{(2)}$  and  $A^{(2)} = L^{(1)^T} E^{(1)} L^{(2)}$ . Terms of the form  $\mathbf{x}^T L^{(1)^T} E^{(1)} \mathbf{q}^*$  and  $\mathbf{y}^T L^{(2)^T} E^{(2)} \mathbf{p}^*$  vanish for Nash equilibria in the interior of strategy space, and we



**Figure 8.** (Colour on-line) Left: Ternary (or barycentric) plot of the components of player 1’s strategy in a realization of the lagging anchor algorithm, for the rock-paper-scissors game. The parameters have the values  $\kappa_i = 0.005$ ,  $\mu_i = \nu_i = 0.05$ ,  $\phi_i = 0.5$ , for which the deterministic dynamics converges. Right: Power spectrum of a component of player 1’s strategy. Solid line: theory, markers: simulation (averaged over 1000 simulations).

will only consider such cases here. The expressions in (10) allow the payoff gradients required for the gradient ascent algorithm to be written in a simple form,

$$\begin{aligned}\nabla_{\mathbf{x}} u_1 &= A^{(1)} \mathbf{y}, \\ \nabla_{\mathbf{y}} u_2 &= A^{(2)} \mathbf{x},\end{aligned}$$

so that the full lagging anchor system (3) becomes

$$\begin{aligned}\mathbf{x}(t+1) &= \mathbf{x}(t) + A^{(1)} \tilde{\mathbf{y}}(t) + \mu_1 (\bar{\mathbf{x}}(t) - \mathbf{x}(t)) \\ \mathbf{y}(t+1) &= \mathbf{y}(t) + A^{(2)} \tilde{\mathbf{x}}(t) + \mu_2 (\bar{\mathbf{y}}(t) - \mathbf{y}(t)) \\ \bar{\mathbf{x}}(t+1) &= \bar{\mathbf{x}}(t) + \nu_1 (\mathbf{x}(t) - \bar{\mathbf{x}}(t)) \\ \bar{\mathbf{y}}(t+1) &= \bar{\mathbf{y}}(t) + \nu_2 (\mathbf{y}(t) - \bar{\mathbf{y}}(t)) \\ \tilde{\mathbf{x}}(t+1) &= \tilde{\mathbf{x}}(t) + \phi_1 (\mathbf{X}(t) - \tilde{\mathbf{x}}(t)) \\ \tilde{\mathbf{y}}(t+1) &= \tilde{\mathbf{y}}(t) + \phi_2 (\mathbf{Y}(t) - \tilde{\mathbf{y}}(t)).\end{aligned}$$

Now, instead of intervals, the players’ strategies are confined to standard simplexes—that is, sets of vectors whose elements are nonzero and sum to one. If the update rules take the strategies outside their simplexes, they must be mapped back to the nearest point in the simplex—this is a ‘convex projection’, described in detail by Michelot [33].

The calculation of the approximations of the covariances and power spectra follow in a similar manner to those for a two-action game. As an example, a phase plot of the lagging anchor dynamics for the ‘rock-paper-scissors’ game with the same values of the constant parameters as that used for the matching pennies game in figure 2 is shown in figure 8, with a corresponding power spectrum. The dynamics is closely comparable to those seen for matching pennies, and the approximation of the power spectrum is similarly accurate.

## 5. Conclusions

Intrinsic noise appears in the lagging anchor dynamics when the players do not know their opponents' strategies. The properties of these fluctuations can be calculated analytically to good approximation, when the deterministic dynamics is stable. In this case, the noise can have a significant effect on the dynamics, inducing quasicycles, similar to those driven by demographic noise in population-based models [28, 29, 30]. See also [22, 23, 24, 25] for studies of intrinsic noise in game learning. In our work we have considered two-player learning in the context of the matching pennies game, and the well-known rock-paper-scissors game. We are able to predict the magnitude of stochastic fluctuations, and to determine the spectral properties of quasicycles analytically, in good agreement with simulations. When the players are asymmetric, i.e. when they use different parameters for their respective lagging anchor adaptation, one of them can take advantage of the resulting quasicycles, achieving a positive average payoff. While we limit the present work to an analysis of the above two-player games, the analytic framework developed here is very general, and can be applied to any game learning system in which the strategies tend to stay close to an equilibrium point. Convergence to Nash points is a key objective with which many adaptation mechanisms in machine learning have been designed, and as we show in our work intrinsic noise due to imperfect sampling can seriously affect whether or not this objective is achieved. The formalism presented in this paper provides a systematic approach with which to estimate deviations from convergence, and we expect that it can be applied to a large class of machine-learning algorithms, beyond the lagging anchor scheme we have discussed here.

## Acknowledgments

JBTS is supported by an EPSRC studentship. TG is grateful for funding by the Research Councils UK (RCUK reference EP/E500048/1), JLS acknowledges funding by EPSRC (grant EP/E050441/1).

## References

- [1] von Neumann J and Morgenstern O 1944 *Theory of Games and Economic Behavior* (Princeton: Princeton University Press)
- [2] Nash J F 1950 *Proc. Natl. Acad. Sci. U.S.A.* **36** 48–9
- [3] Nash J F 1951 *Annals of Mathematics* **54** 286–95
- [4] Sato Y, Akiyama E and Farmer J D 2002 *Proc. Natl. Acad. Sci. U.S.A.* **99** 4748–51
- [5] Sato Y and Crutchfield J P 2003 *Phys. Rev. E* **67** 015206
- [6] Camerer C F and Ho T H 1999 *Econometrica* **64** 827–74
- [7] Ho T H, Camerer C F and Chong J K 2007 *Journal of Economic Theory* **133** 177–98
- [8] Macy M W and Flache A 2002 *Proc. Natl. Acad. Sci. U.S.A.* **99** 7229–36
- [9] Young P H 2004 *Strategic Learning and its Limits* (Oxford: Oxford University Press)
- [10] Fudenberg D and Levine D K 1998 *The Theory of Learning in Games* (Cambridge: MIT Press)
- [11] Alpaydin E 2004 *Introduction to Machine Learning* (Cambridge: MIT Press)
- [12] Dahl F A 2005 *Journal of Economic Behaviour and Organization* **57** 287–303

- [13] Gintis H 2000 *Game Theory Evolving* (Princeton: Princeton University Press)
- [14] Sutton R S and Barto A G 1998 *Reinforcement Learning: an Introduction* (Cambridge: MIT Press)
- [15] Chong E K P and Źak S H 2008 *An Introduction to Optimization* (Wiley-Interscience Series in Discrete Mathematics and Optimization vol 72) (Hoboken: Wiley-Interscience)
- [16] Singh S, Kearns M and Mansour, Y 2000 *Proc. 16th Conf. on Uncertainty in Artificial Intelligence (Stanford)* (San Francisco: Morgan Kaufmann) p 541
- [17] Butterworth J M and Shapiro J L 2009 *Proc. 10th Workshop on Foundations of Genetic Algorithms (Orlando)* (New York: ACM)
- [18] Selten R 1991 *Game Equilibrium Models* vol 1 *Evolution and Game Dynamics* ed R Selten (Berlin: Springer-Verlag)
- [19] Crawford V P 1974 *Econometrica* **42** 885–91
- [20] Dahl F A 2001 *12th Eur. Conf. on Machine Learning (Freiburg)* ed L de Raedt and P Flach (Berlin: Springer) p 85
- [21] Dahl F A 2002 *Machine Learning* **49** 5–37
- [22] Galla T 2009 *Phys. Rev. Lett.* **103** 198702
- [23] Galla T 2011 *J. Stat. Mech. Theor. Exp.* P08007
- [24] Bladon A J and Galla T 2011 Learning to play public good games *Preprint* arXiv:1107.0878
- [25] Realpe-Gomez J, Szczesny B, Dall'Asta L and Galla T 2011 Fixation and escape times in stochastic game learning *Preprint* arXiv:1102.0876
- [26] Kitagawa G 1977 *International Journal of Control* **25** 745–53
- [27] Michel A N, Hou L and Liu D 2008 *Stability of Dynamical Systems: Continuous, Discontinuous, and Discrete Systems* (Berlin: Birkhäuser)
- [28] McKane A J and Newman T J 2005 *Phys. Rev. Lett.* **94** 218102
- [29] Reichenbach T, Mobilia M and Frey E 2006 *Phys. Rev. E* **74** 051907
- [30] Bladon A J, Galla T and McKane A J 2010 *Phys. Rev. E* **81** 066122
- [31] Traulsen A, Röhl T, and Schuster H G 2004 *Phys. Rev. Lett.* **93** 028701
- [32] Röhl T, Traulsen A, Claussen J C and Schuster H G 2008 *Phys. Rev. E* **78** 026108
- [33] Michelot C 1986 *Journal of Optimization Theory and Applications* **50** 195–200